VECTRA®
SECURITY THAT THINKS

Attacker Behavior Industry Report

VECTRA®
SECURITY THAT THINKS

TABLE OF CONTENTS

The Vectra® Attacker Behavior Industry Report provides a first-hand analysis of active and persistent attacker behaviors inside the enterprise networks of Vectra customers.

The report examines a wide range of cyberattack detections and trends from a sample of 350 Vectra Cognito® deployments with more than 5 million hosts per month from nine different industries. This report takes a multidisciplinary approach that spans all strategic phases of the attack lifecycle.

While there are plenty of threat-research reports, this one offers unique insights about real-world cyberattacker behaviors found in cloud, data center and enterprise networks.

Most security industry reports focus on statistics of known threats, such as exploits and malware families, or provide a post-mortem of successful breaches. The Vectra Attacker Behavior Industry Report takes a multidisciplinary approach that spans all strategic phases of the attack lifecycle. It presents data by specific industries that highlight relevant differences between them.

Dataset and demographics

The data in this report is based on anonymized network metadata from 350 Vectra Cognito deployments that have opted to share detection metrics. The Cognito platform identifies behaviors that indicate attacks in progress by directly monitoring all network traffic and relevant logs, including traffic to and from the internet, internal traffic between network devices, and virtualized workloads in private data centers and public clouds.

This analysis provides important visibility into advanced phases of attacks. The Cognito platform detects threats that bypass perimeter security controls and observes the progression of the attack after an initial compromise.

Observations

- Across all industries, in the six-month period from July to December 2019, there was an average of 215 attacker behavior detections per 10,000 hosts with a peak of 252 in July 2019.
- Technology (138 detections per 10,000) and education organizations (102 detections per 10,000) remain the most common sectors to exhibit command & control behaviors. Everyone else experienced just 39 detections per 10,000 hosts.
- It is rare to see large volumes of TOR traffic in any organization as it serves few if any legitimate business purposes. Across all industries TOR averaged 3 detections per 10,000 hosts. In December, it averaged 19 detections per 10,000 hosts driven by a spike in December in technology companies in the Asia-Pacific region.
- Finance and insurance organizations experienced 29 port scan detections per 10,000 hosts, compared to an industry average of 11.
- Government agencies detected the lowest rate of reconnaissance behaviors, at 32 per 10,000 hosts, while finance and insurance organizations detected the highest at 93 per 10,000 hosts.
- File server (SMB) brute force behaviors attempting to crack user passwords were observed a year-high 22 times per 10,000 hosts in July. Over the rest of the year (August-December), SMB brute force was observed 14 times per 10,000 hosts.
- Small companies (1 to 3,000 employees) are more at risk of lateral movement attacks. Small companies observed 112 lateral movement behaviors per 10,000 hosts while medium and large companies detected 64.
- Vectra customers achieved a 38X workload reduction for Tier-1 analysts by automating the process of detection, triage, correlation and prioritization of security incidents to hosts, enabling security operations teams to focus on compromised hosts that pose the highest risk.

Some key terms used in this report

Conventional terminology is used where possible in this report, but some things are not clear or a universally accepted term. In such cases, a term related to how the Vectra Cognito platform describes data is used consistently. This section should provide clarity to those terms.

Hosts observed – unique instance of a host based on a variety of artifacts. Cognito tracks every host independent of IP address by producing a unique identifier for each host. A host can be a physical device or virtual/cloud workload.

Events flagged – individual alerts that make up a detection. Events are a specific instance of the occurrence of a behavior, such as a single port scan or remote connection activity.

Detection – correlation of individual events to a single instance that creates a specific attacker behavior. Detections are based on machine learning algorithmic models and labeled according to the model that produced the detection.

Hosts with detection – all detection events are correlated to specific hosts that show signs of threat behaviors. Each host is scored based on the combination of attacker behavior detections that cover multiple phases of the attack lifecycle.

The Attacker Behavior Industry Report also presents data by specific industries and highlights relevant differences between industries.

Figure 1 and 2 provide demographic information on the organizations included in this analysis. The first offers two views of the industries represented, one based on the percent of organizations and the other based on the total number of detections within a combined six-month period.

Figure 2 shows the distribution across small, medium, and large organizations. It also gives a range of employee counts for each of those categories to make it easy to determine where an organization would fit. (Large = 25001+, medium = 5001-25k, small = 0-5k)

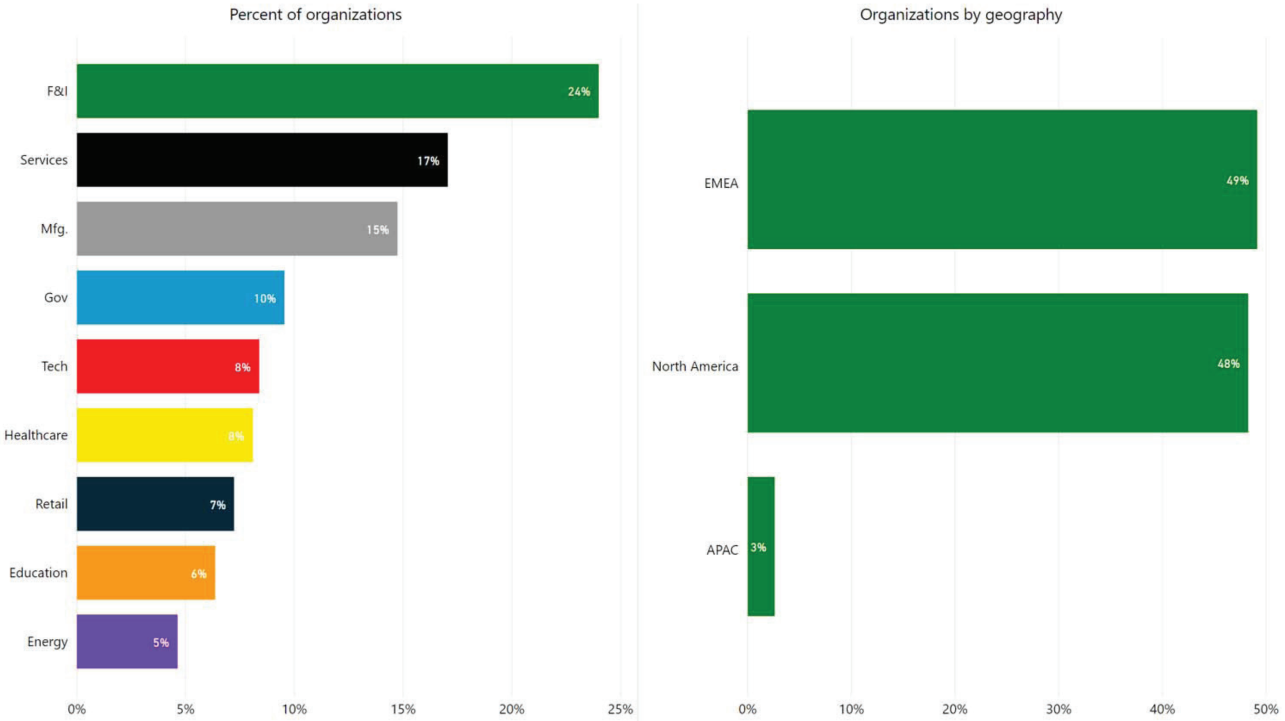

Figure 1: Industry demographics by percent of number of organizations, total number of events flagged and geography

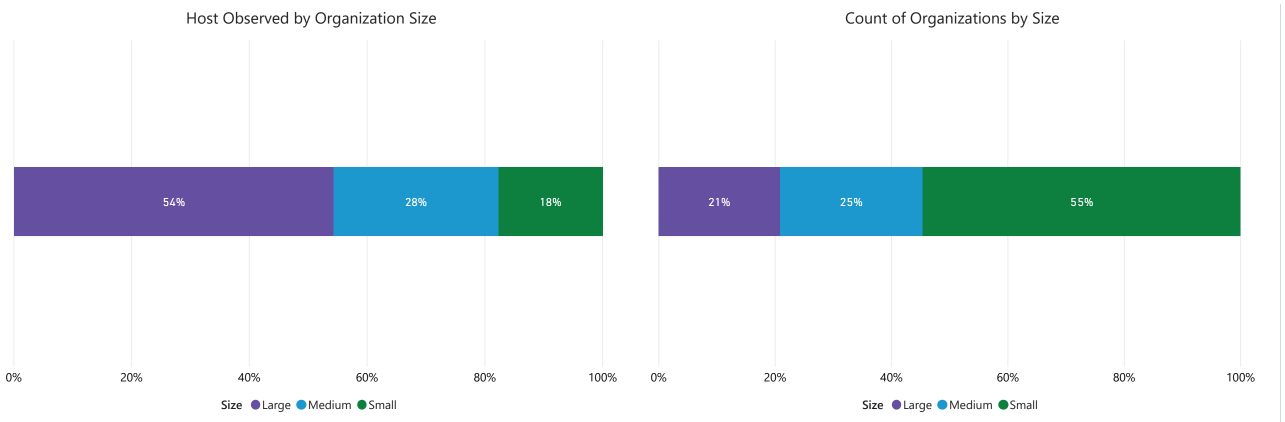

Figure 2: Industry demographics by organization size

Operational efficiency and ROI

Cybersecurity is an ongoing exercise in operational efficiency. Organizations have limited resources to address unlimited risks, threats and attackers. This means that security products must always be evaluated in terms of efficiency and their impact on the operational fitness of the organization.

Time is the most important factor in detecting network breaches. To mitigate damage, attacks must be detected in real time before key assets are stolen or damaged. Traditionally, detecting and responding to targeted attacks is a very time-consuming process and requires security teams to manually sort through mountains of alerts.

Cognito Detect™ from Vectra continuously detects and prioritizes threats at speed using AI. These attacker behaviors are correlated with compromised devices, which are in turn correlated with common attack vectors and larger attack campaigns. Thousands of threat indicators are reduced to hundreds of attacker behaviors on dozens of devices that can be part of broader attack campaigns.

There was a wide variance in the size of the networks analyzed, with the smallest consisting of a few hundred devices and workloads to the largest networks with more than 400,000.

To account for this variance, the data has been normalized to a network with 10,000 devices and workloads, making it easier to compare the prevalence of threats in a network on a per capita basis. Any device with an IP address – including IoT devices, smartphones, tablets and laptops – are monitored in addition to servers and virtual workloads.

Figure 3 provides a normalized view of detections per 10,000 by industry over the six-month period, from July 2019 – December 2019, of this report.

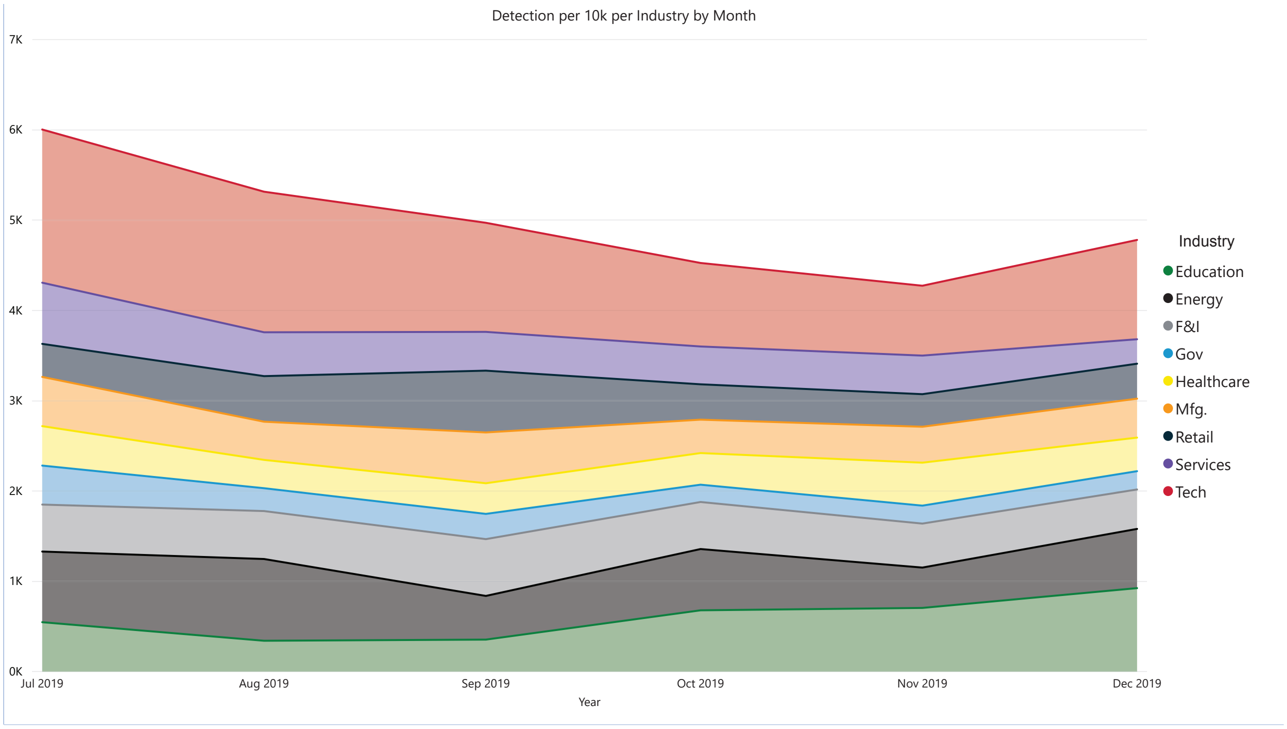

Figure 3

Overall, Cognito reduced the investigation workload of security analysts by 38X, compared to manually investigating all attacker behaviors and compromised host devices.

Reduction in workload for Tier-1 security analysts

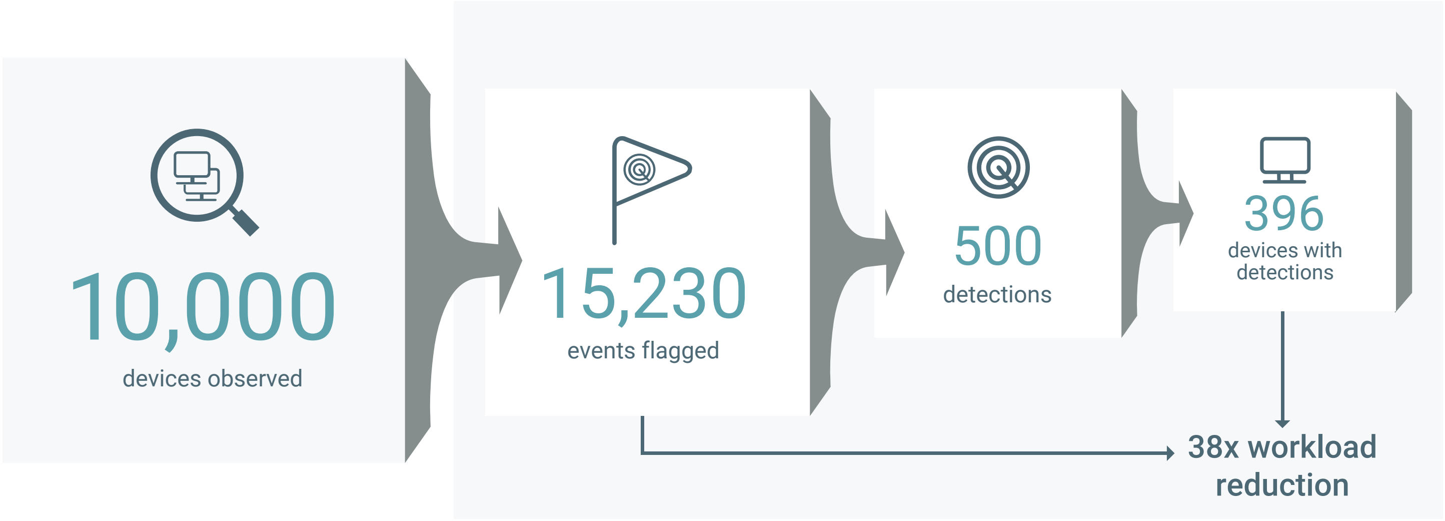

10,000
devices observed

15,230
events flagged

500
detections

396
devices with detections

38x workload reduction

REDUCTION IN WORKLOAD PER 10,000 DEVICES OBSERVED BY INDUSTRY

Industry	Events flagged	Detections	Devices with detections	Critical severity	High severity	Workload reduction
Education	12,666	617	650	11	34	19x
Energy	18,617	598	338	8	16	55x
F&I	19,510	511	361	8	20	54x
Government	9,328	249	268	6	12	35x
Healthcare	15,423	381	411	7	16	38x
Manufacturing	17,064	431	323	6	15	53x
Retail	9,437	421	283	7	24	33x
Services	14,679	430	365	12	20	40x
Tech	18,100	1,193	634	12	23	29x

Scoring

Cognito from Vectra monitors individual devices and workloads for extended periods of time and attributes detections to any device or workload that behaves suspiciously. The detection scores and when they occurred are key inputs for the host device scores.

Cognito's scoring is comprised of two dynamic metrics – threat and certainty scores – applied to individual detections and the host devices against which they are reported.

The threat score of a detection expresses the potential for harm if the security event is true (for example, if spamming behavior or data exfiltration was occurring). Because a threat is a measure of the potential for harm, it reflects worst-case scenarios.

The certainty score of a detection reflects the probability that a given security event occurred (for example, the probability of spamming behavior occurring or the probability of data exfiltration occurring), given all the evidence observed so far.

Certainty is based on the degree of difference between the threat behavior that caused the detection and normal behavior. As such, the certainty score of an individual detection changes over time.

Since detections are dynamic, changes in their scores cause changes to attributed host device scores. Critical and high scores help security analysts prioritize their investigation efforts because they represent behaviors with the highest certainty and greatest potential to cause significant damage.

Other factors that influence host device scores include repetition of an observed detection or a combination of detections that indicate a cyberattack is progressing toward its objective.

Every detection type has a maximum lifespan, ranging from a few days to a month. When a detection has no recurring activity, its effect on a host device score will slowly decline to zero. A detection past its maximum lifespan becomes inactive and has no impact on the host device score.

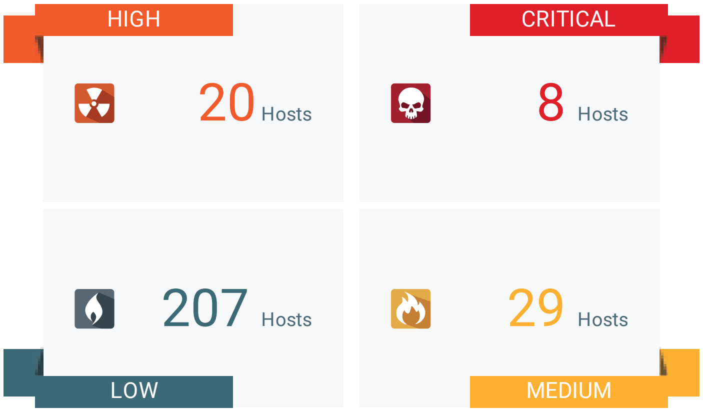

An overview of detections per 10,000 devices and workloads

On average, for every 10,000 devices and workloads monitored in a one-month period, the peak count of host severity was 8 critical and 20 high. These devices and workloads present the greatest threat to the organization and require a security analyst's immediate attention.

When breaking down host device severity statistics across vertical industries, Vectra benchmarked the volume of devices and workloads prioritized for each severity in relation to each vertical and with the overall average.

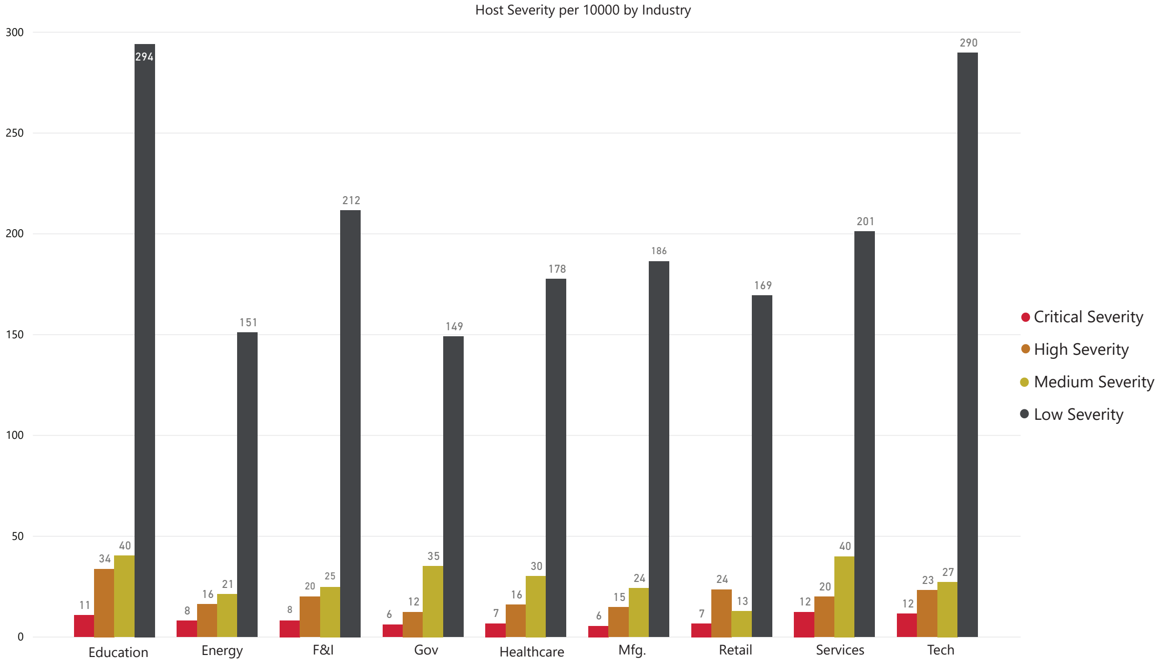

For example, the number of low alerts in technology and higher education organizations are 50% higher than the normal rate, which is indicative of attacker behaviors early in the attack lifecycle, but do not necessarily represent the more critical steps of an attack like data theft.

Inversely, the government agencies have a low volume of hosts prioritized as high or critical, which indicates a low rate of behaviors progressing deep across the attack lifecycle.

Threats by type per 10,000 devices

To dig deeper, Vectra provides a breakdown of detection statistics by industry. The charts below show threat behaviors across the attack lifecycle. These behaviors are strong indicators of exposure and risk in an organization and enable security analysts to focus their time and effort on what matters most.

While not every stage is necessary in an attack, they are interrelated, and we often see an attack progress through the stages with the ultimate outcome of financial gain, data exfiltration or data destruction.

This data represents in-progress attacker behaviors. Activity like command and control (C&C) and reconnaissance occur in the earlier stages of an attack, enabling organizations to quickly mitigate the threat before it can spread. These are the most common detected behaviors.

Behaviors like lateral movement occur later in the attack lifecycle as cybercriminals strengthen their foothold in an organization by stealing administrative credentials to access servers. These types of detections warrant high-priority action from incident response teams to prevent irreversible damage from a data exfiltration.

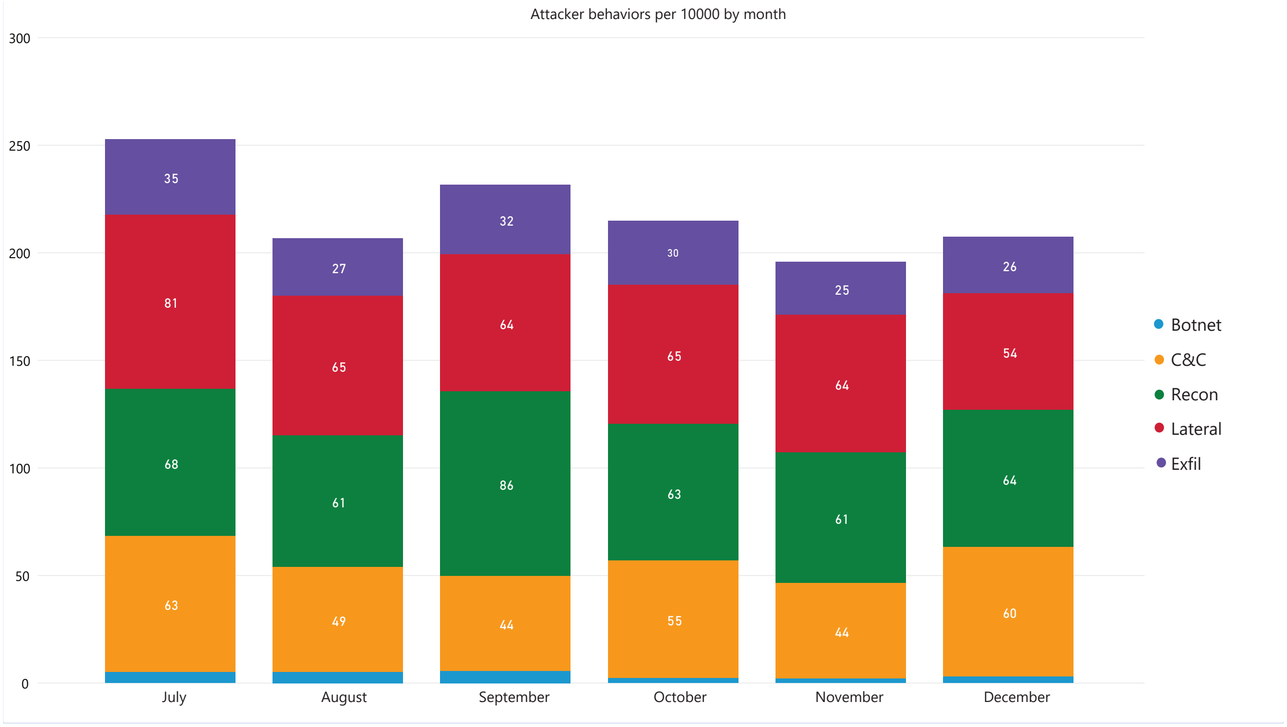

Botnets

Botnets are opportunistic attack behaviors in which a device makes money for its bot herder. The ways in which an infected device can be used to produce value can range from mining bitcoins to sending spam emails to producing fake ad clicks. To turn a profit, the bot herder utilizes devices, their network connections and, most of all, the unsullied reputation of their assigned IP addresses.

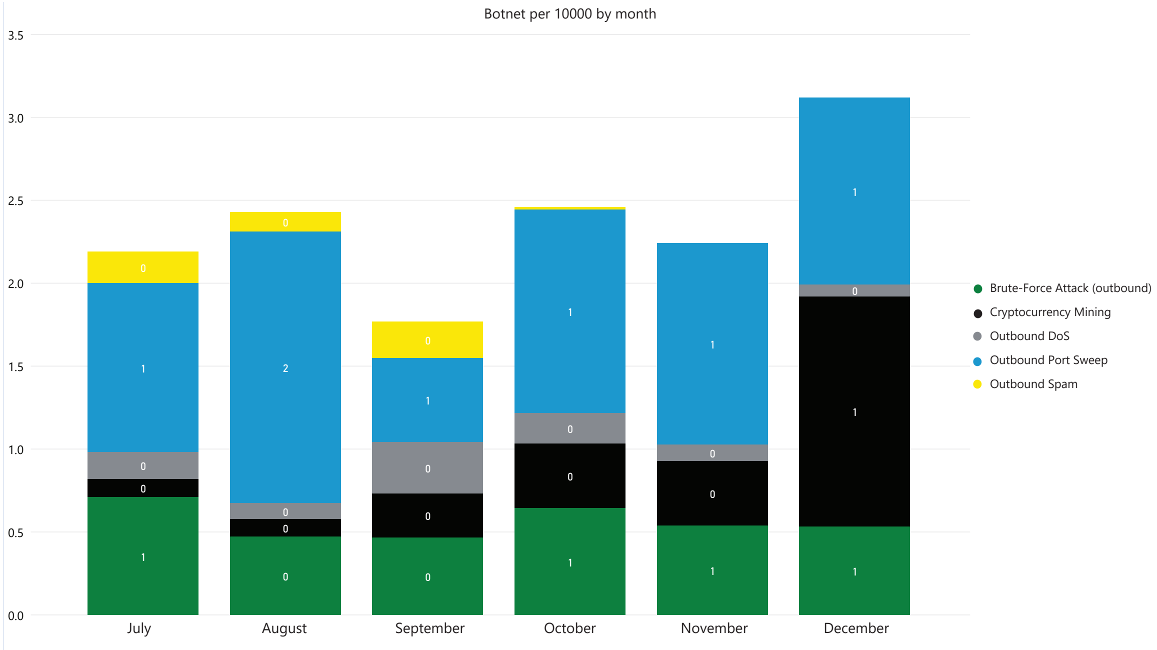

Command and control

C&C traffic occurs when a device appears to be under control of an external malicious entity. Most often, the control is automated because the device is part of a botnet or has adware or spyware installed.

Rarely, but most importantly, a device can be manually controlled by a nefarious outsider. This is the most threatening case and it often means the attack is targeted at a specific organization.

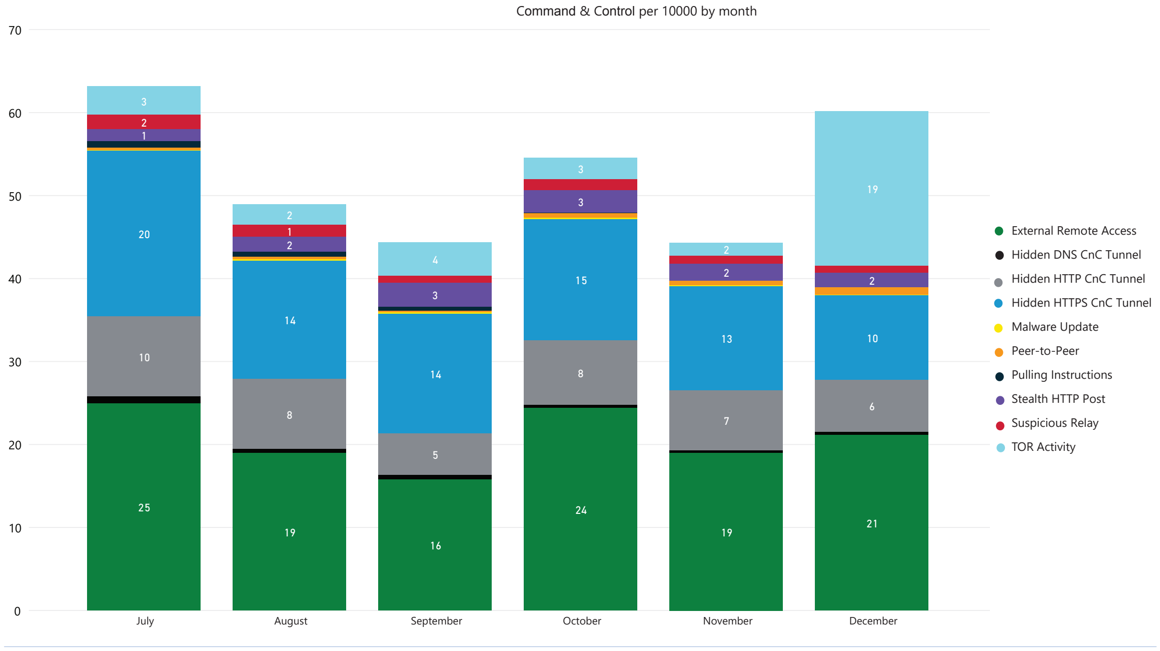

Reconnaissance

Reconnaissance attacker behaviors occur when a device is used to map-out the enterprise infrastructure. This activity is often part of a targeted attack, although it might indicate that botnets are attempting to spread internally to other devices. Detection types cover fast scans and slow scans of systems, network ports and user accounts.

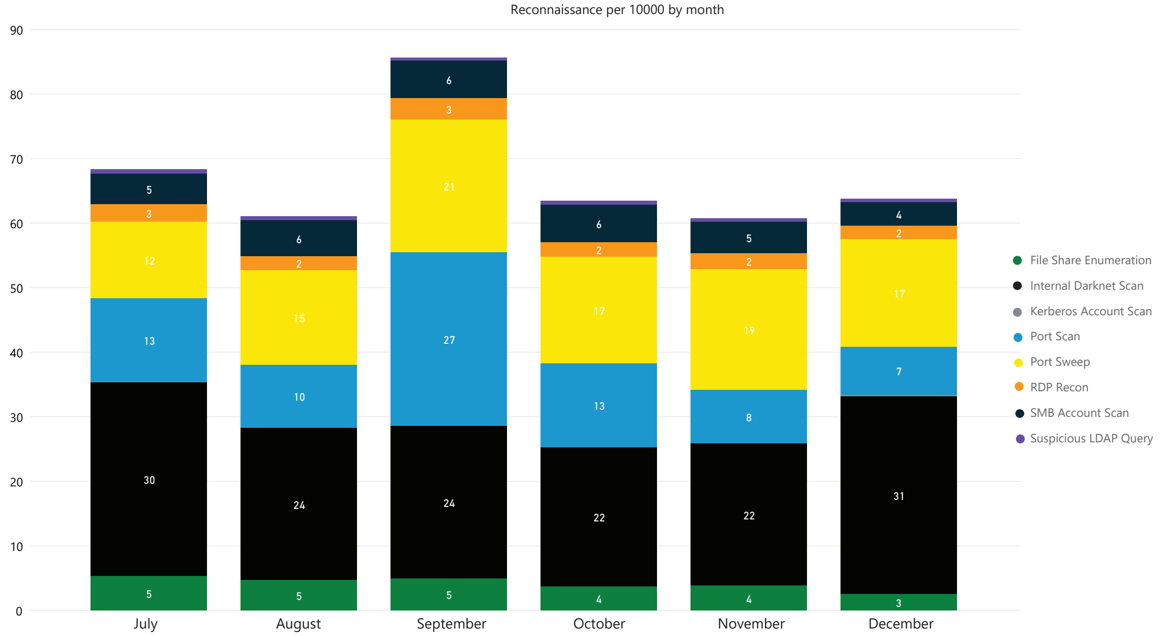

Lateral movement

Lateral movement covers scenarios of lateral action meant to further a targeted attack. This can involve attempts to steal account credentials or to steal data from another device.

It can also involve compromising another device to make the attacker's foothold more durable or to get closer to target data. This stage of the attack lifecycle is the precursor to moving into private data centers and public clouds.

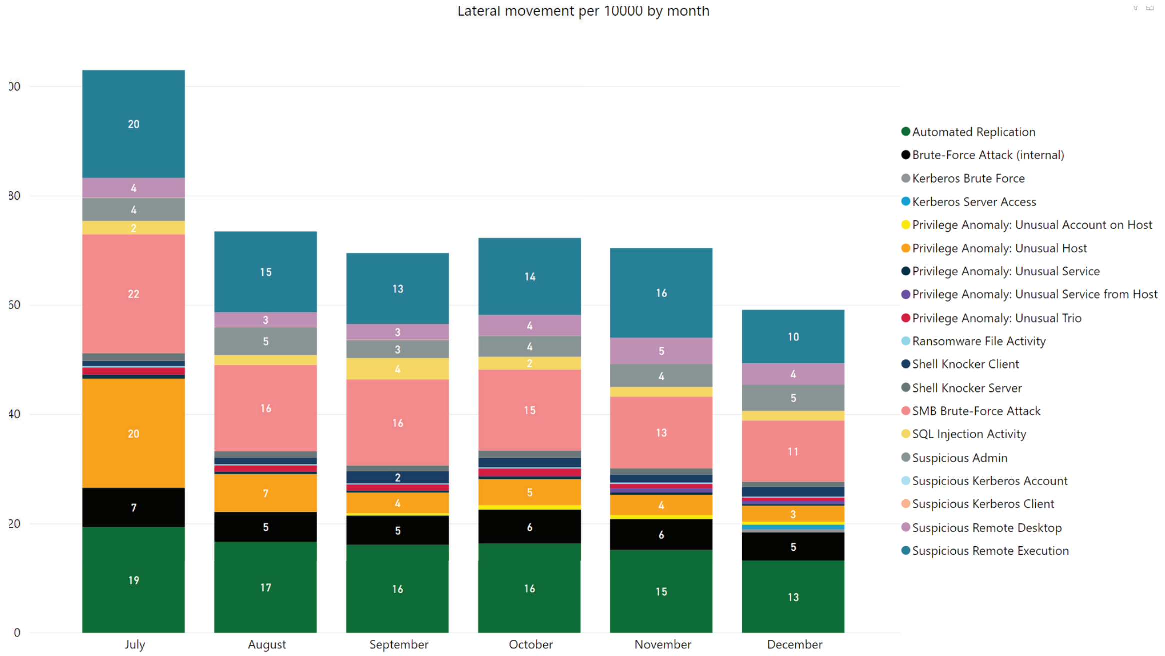

Data exfiltration

Data exfiltration behaviors occur when data is sent to the outside in a way that is meant to hide the transfer. Normally, legitimate data transfers do not involve the use of techniques meant to hide the transfer. The device transmitting the data, where it is transmitting the data, the amount of data and the technique used to send it are indicators of exfiltration.

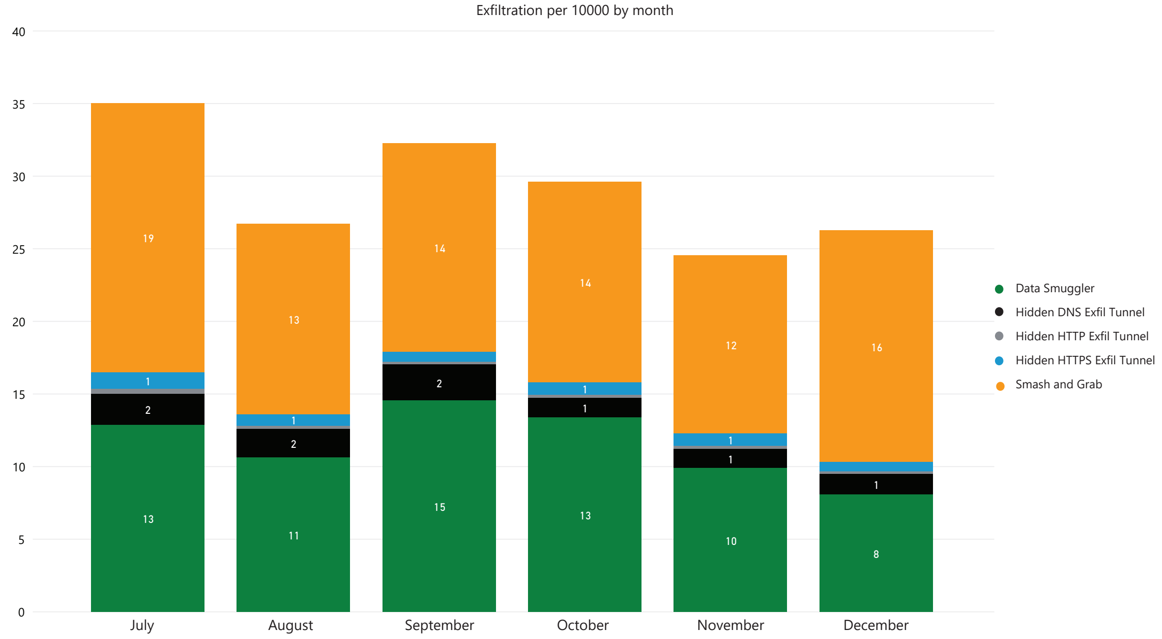

Threats by industry per 10,000 devices

The chart below shows the volume of threat detections that were triggered in each industry. This view shows how each industry fared per capita as well as which industries generated the most detections by volume.

The technology industry and education represent the highest percentage of detections across all industries, primarily due to a high volume of C&C (education) and reconnaissance (technology).

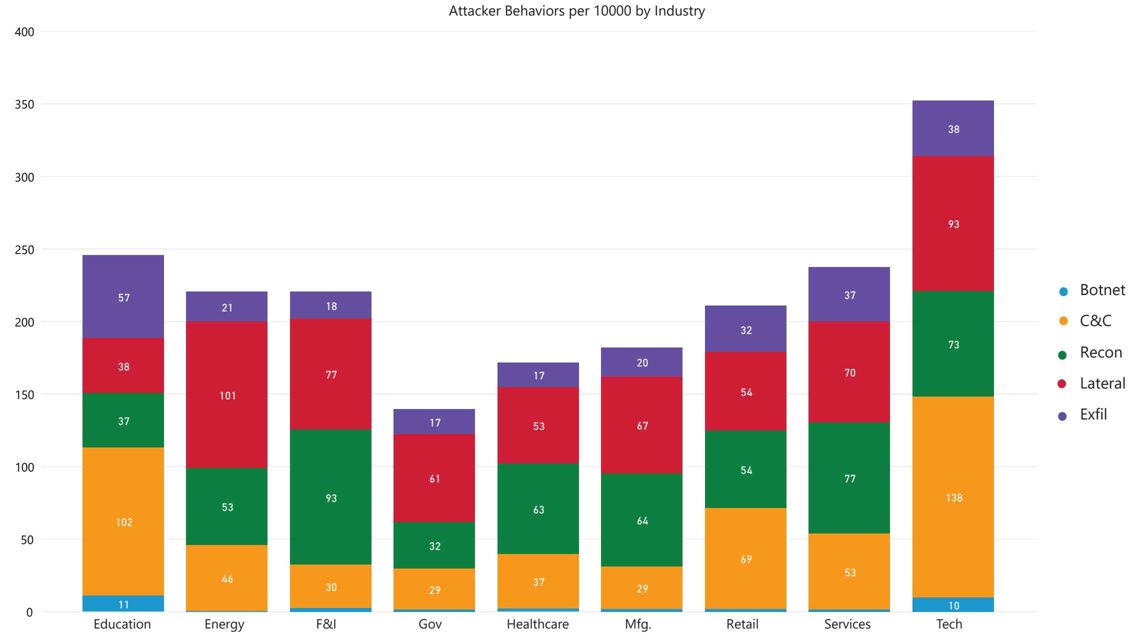

Botnets by industry

Education has and will always be a hotbed for opportunistic botnet behaviors. This is due to the open nature of universities and the curiosity of students. They have always been a place where cryptomining and abnormal ad and web activity have thrived.

In the second half of 2019 there was a surge in activity within technology companies. While the volume of botnet activity is still low relative to the overall landscape of hosts and behaviors, there is a noticeable uptick in cryptomining and outbound port sweep activity within these technology organizations.

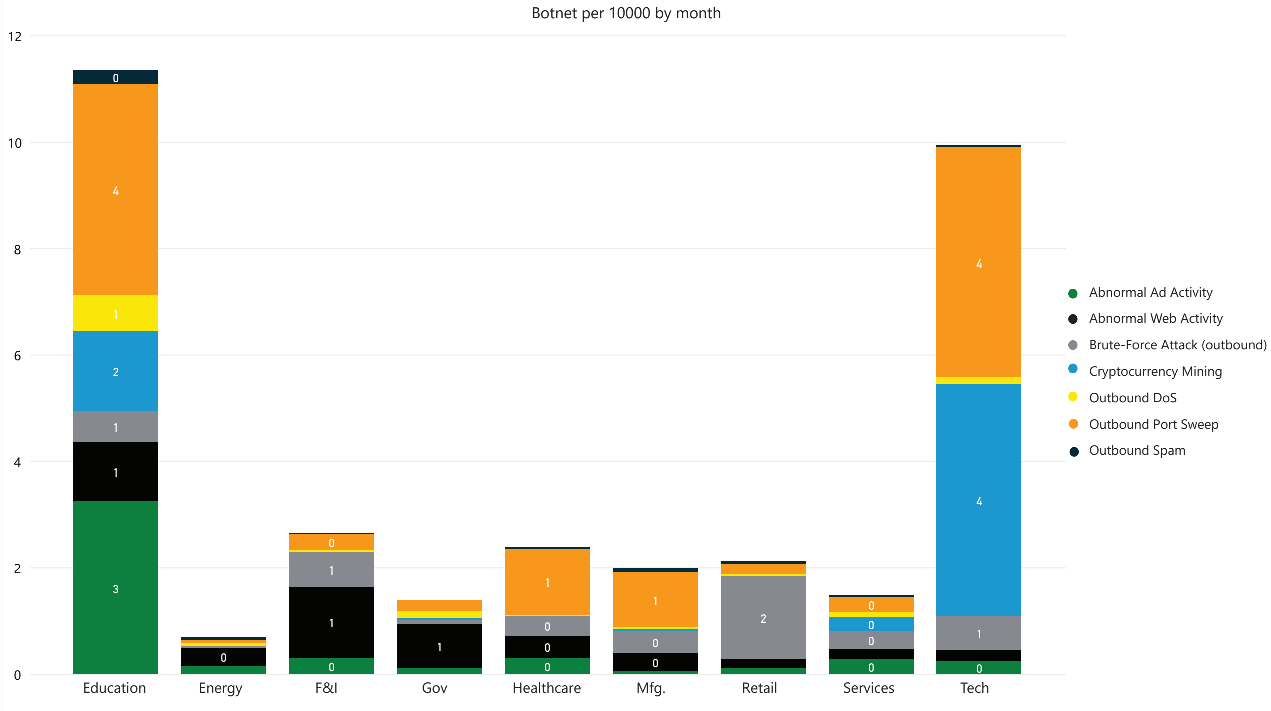

C&C by industry

The technology industry has set a trend in 2019 of exhibiting a large amount of external remote access and hidden HTTPS tunnels. Joining technology in the last six months in this trend are education and retail organizations. In particular with a high rate of external remote access.

External remote access is a common occurrence in every industry as remote workers are an essential part of the workforce. These remote workers use remote desktop access software to access hosts from the outside (e.g. GotoMyPC, RDP). Attackers use the same tactics to control a host using malware with remote access capability (e.g. Meterpreter, Poison Ivy) that connects to its C&C server and receives commands from a human operator.

The use of a hidden tunnel by some software may be benign, but it represents significant risk as the intention is to bypass security controls. Hidden tunnels used as part of a targeted attack are meant to slip by perimeter security controls and indicate a sophisticated attacker.

What is new for the second half of 2019 is a spike in TOR traffic. It is rare to see large volumes of TOR traffic in any organization as it serves few if any legitimate business purposes. Across all industries TOR averaged 3 detections per 10,000 hosts. In December, it averaged 19 detections per 10,000 hosts driven by a spike in December in technology companies in the Asia-Pacific region.

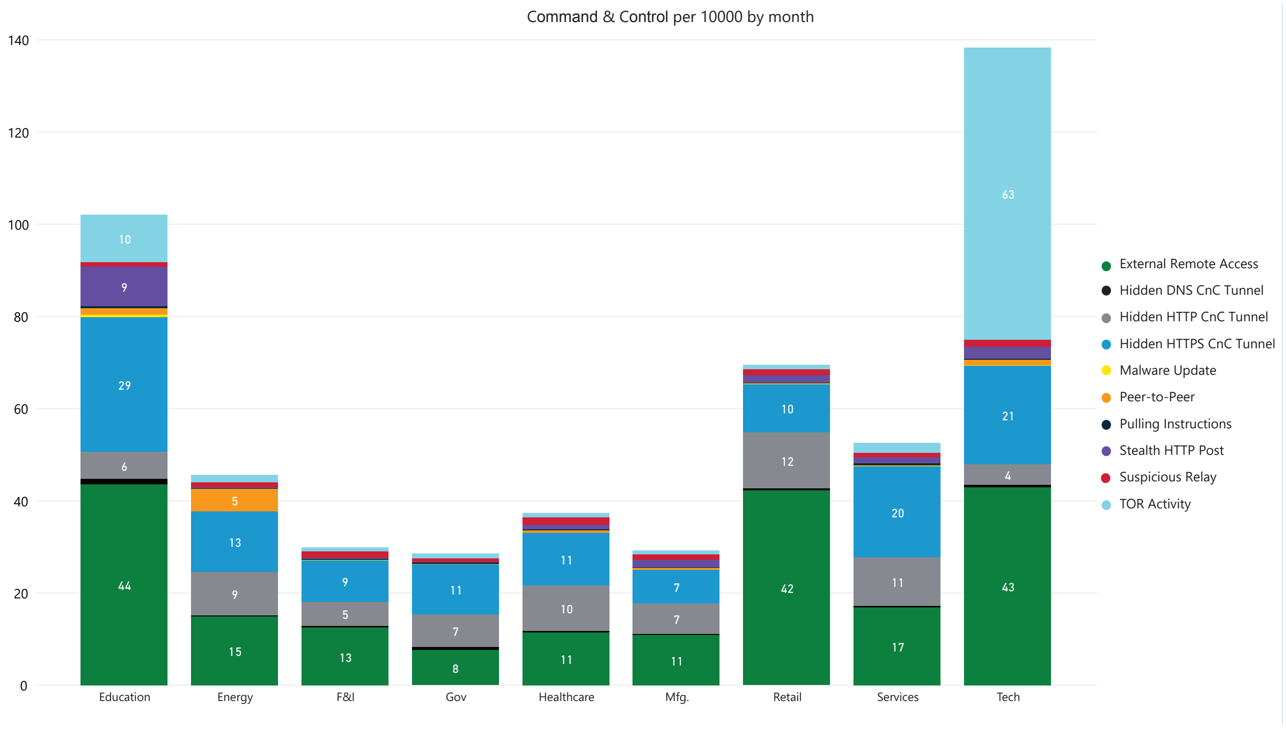

Reconnaissance by industry

In the world of cyberattacks, doing simple stuff always beats trying complicated techniques. We see that trend in reconnaissance with finance and insurance organizations experiencing 29 port sweep detections per 10,000 hosts, compared to an industry average of 11. Port sweeps are easy to do, necessary to map networks for further lateral movement, and most importantly, they work.

As part of a port sweep, scans of nonexistent IP addresses on the network (internal darknet scans) are the most common reconnaissance behavior. Internal darknet scans represent asset-mapping systems in the network or hosts that have moved to a new network and are unsuccessfully attempting to connect to many previously available services. While the bulk of darknet scans occurs from internal systems like vulnerability scanners and asset discovery systems, slow reconnaissance of systems may represent the beginning of a targeted attack in the network.

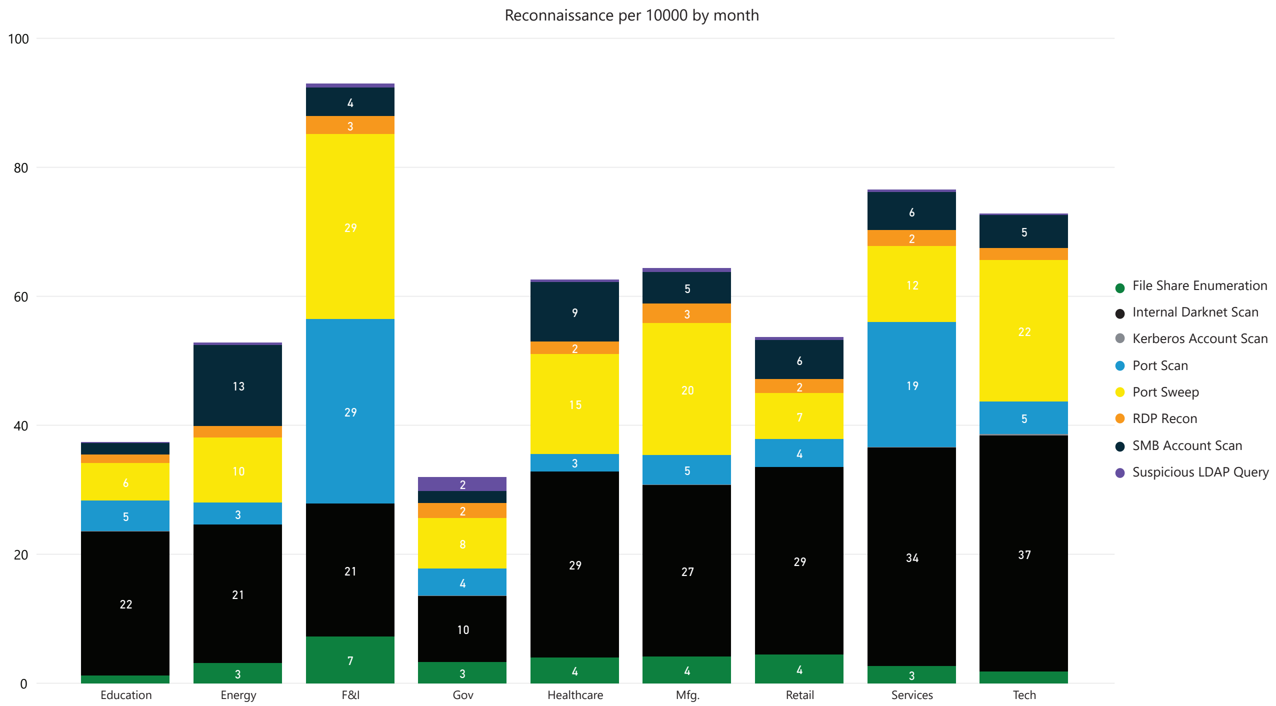

Lateral movement by industry

Again, noisy attack techniques are a standout across all industries. SMB brute force behaviors were observed a year-high 22 times per 10,000 hosts in July. Over the rest of the year (August-December), SMB brute force behaviors only occurred 14 times per 10,000 hosts.

SMB brute force behaviors indicate that a device is making multiple login attempts, using the same accounts, to access a file server. Successful harvesting of account credentials (usernames and password) of other accounts, particularly more privileged accounts, is a classic progression of a targeted attack. Even if trigged due to a misconfiguration, the identified behavior is creating significant stress on the target system and should be cleaned up.

On another note, small companies observed 112 lateral movement behaviors per 10,000 hosts while medium and large companies detected 64. Does this mean it is easier to move around in a small shop than a large one? Perhaps the large number of ransomware attacks in 2019 targeting these smaller organizations is a good indicator of success of how easily attackers can move from point a to b without being noticed.

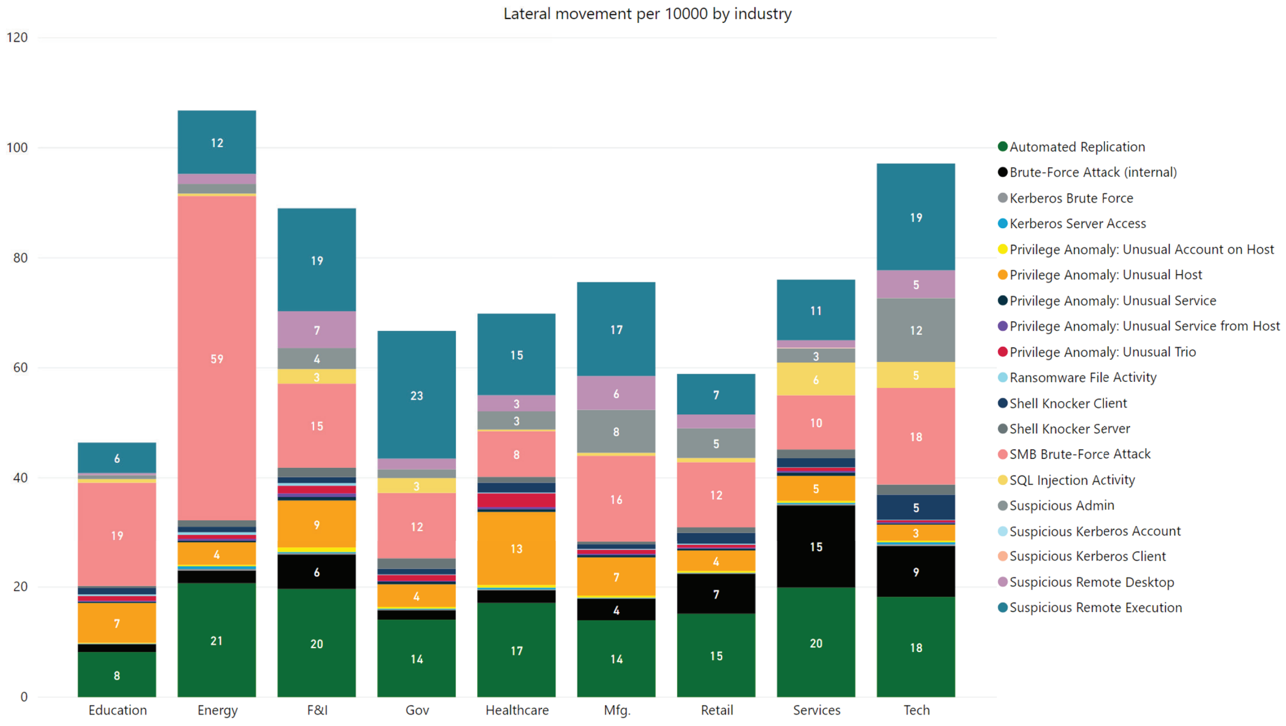

Exfiltration by industry

Smash-and-grab is the most common exfiltration behavior across all industries. It is triggered when a device transmits unusually large volumes of data to destinations that are not considered normal for the environment within a short amount of time. If the external service to which data was uploaded is not an IT-sanctioned service, the potential business risk is high. We saw a significant uptick in exfiltration behaviors in education organizations, with smash-and-grab behaviors leading the way.

The second most common exfiltration behavior is data smuggler, which was commonly observed across all industry verticals. This behavior is detected when an internal host acquires a large amount of data from one or more servers and sends significant volumes of data to an external system. While acquiring and transmitting a large quantity of data to the outside within a short period of time may be pure coincidence, the outbound data transfer is significant enough to warrant further examination.

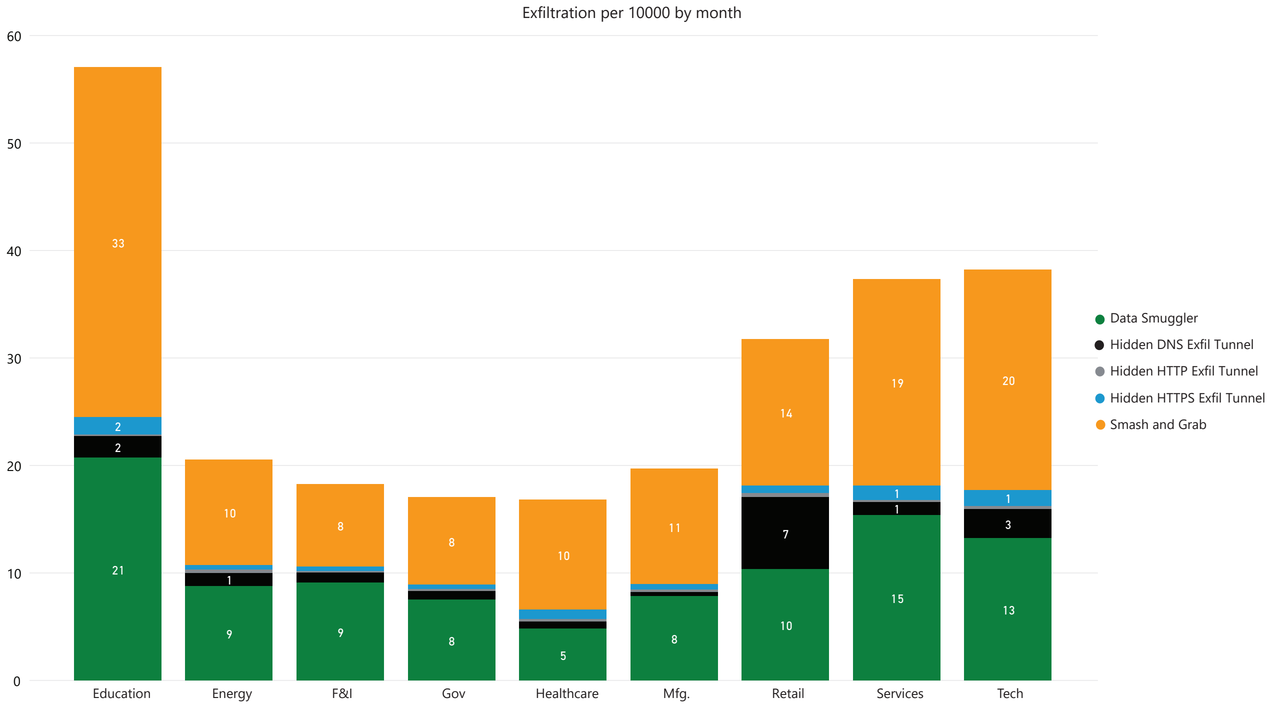

Conclusion

There is a major security gap that's obvious, important and urgent: the ability to know if your own networks are compromised. Prevention will fail, and attackers will get inside. Legacy network security is the weak link.

Comprehensive, enterprise-wide threat-detection and response is mandatory in today's hostile data environments, and the stakes have never been higher. Closing the gap starts with having the right data with the right context to provide continuous awareness of compromise.

Cognito was designed from the ground up to solve this problem. Cognito delivers high-fidelity network data – knowledge of what's happening in every conversation – enriched with context specific to security applications, such as the names of hosts, existence of beacons, and the privilege level of accounts.

Cognito scales efficiently to even the largest enterprise networks to provide 360-degree visibility across users, data centers and cloud – now and in the past.

Vectra would like to thank the organizations who opted-in to share network metadata that was analyzed for this report.

VECTRA®
SECURITY THAT THINKS

Email info@vectra.ai www.vectra.ai